

# The Kepler End-to-End Data Pipeline: From Photons to Far Away Worlds

Brian Cooke  
Jet Propulsion Laboratory  
California Institute of Technology  
4800 Oak Grove Dr.  
Pasadena, CA 91109  
818-354-4605  
Brian.C.Cooke@jpl.nasa.gov

Richard Thompson  
Orbital Sciences Corporation  
NASA Ames Research Center  
Moffett Field, CA 94035  
650-604-5047  
Richard.S.Thompson@nasa.gov

Shaun Standley  
Jet Propulsion Laboratory  
California Institute of Technology  
4800 Oak Grove Dr.  
Pasadena, CA 91109  
818-393-0936  
Shaun.Standley@jpl.nasa.gov

## TABLE OF CONTENTS

1. KEPLER MISSION OVERVIEW .....	1
2. END TO END DATA PATH.....	4
3. KEPLER OPERATIONS.....	6
4. SUMMARY .....	8
REFERENCES.....	8

**Abstract** — Launched by NASA on 6 March 2009, the Kepler Mission has been observing more than 100,000 targets in a single patch of sky between the constellations Cygnus and Lyra almost continuously for the last two years looking for planetary systems using the transit method. As of October 2011, the Kepler spacecraft has collected and returned to Earth just over 290 GB of data, identifying 1235 planet candidates with 25 of these candidates confirmed as planets via ground observation. Extracting the telltale signature of a planetary system from stellar photometry where valid signal transients can be as small as a 40 ppm is a difficult and exacting task. The end-to-end process of determining planetary candidates from noisy, raw photometric measurements is discussed.

The Kepler mission is described in overview and the Kepler technique for discovering exoplanets is discussed. The design and implementation of the Kepler spacecraft, tracing the data path from photons entering the telescope aperture through raw observation data transmitted to the ground operations team is described. The technical challenges of operating a large aperture photometer with an unprecedented 95 million pixel detector are addressed as well as the onboard technique for processing and reducing the large volume of data produced by the Kepler photometer.

The technique and challenge of day-to-day mission operations that result in a very high percentage of time on target is discussed. This includes the day to day process for monitoring and managing the health of the spacecraft, the annual process for maintaining sun on the solar arrays while still keeping the telescope pointed at the fixed science target, the process for safely but rapidly returning to science operations after a spacecraft initiated safing event and the long term anomaly resolution process.

The ground data processing pipeline, from the point that science data is received on the ground to the presentation of preliminary planetary candidates and supporting data to the science team for further evaluation is discussed. Ground management, control, exchange and storage of Kepler's large and growing data set is discussed as well as the process and techniques for removing noise sources and applying calibrations to intermediate data products.

## 1. KEPLER MISSION OVERVIEW

The Kepler mission began on 6 March 2009 with a spectacular night-time launch aboard a workhorse Delta II 7625-10L launch vehicle from Cape Canaveral, Florida. Launching into an Earth trailing heliocentric orbit, Kepler began a mission to find earth-sized and larger planets circling distant stars. After successful early operations and commissioning, Kepler began to observe a small patch of sky in the Cygnus region nearly continuously for 3½ years. Kepler's single instrument, a 0.95m aperture Schmitt photometer, uses a 95 million pixel CCD array to monitor the brightness of over 100,000 target stars for the miniscule dimming that reveals a planet transiting the line of sight between Kepler and its host star.

Kepler's imaging technique co-adds 270 exposures, each several seconds long, to form a 30 minute cadence. Significant onboard data processing and compression

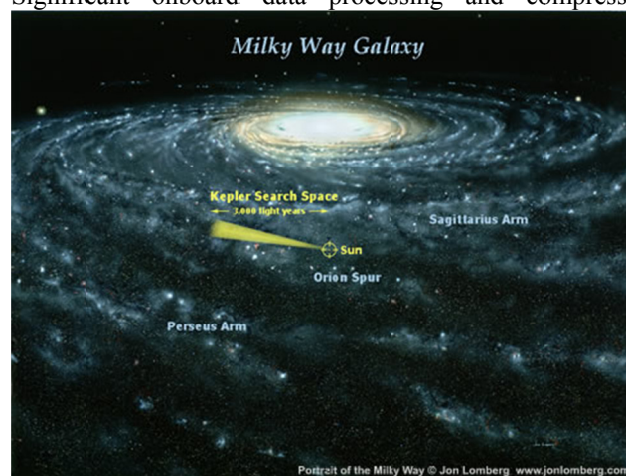


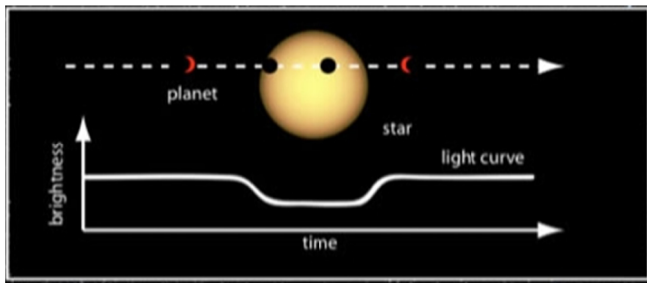
Figure 1 – An illustration of Kepler's search space [4]

reduces the data set to manageable volumes. Once a month, Kepler is turned away from the science field to transmit just over 10 Gb of photometry data over a high bandwidth Ka-band radio link to the ground. Once successfully transmitted, Kepler turns back to the science field to begin another month's observation. Every third month, Kepler is rotated 90° about the spacecraft / science field vector to keep the body fixed solar arrays illuminated.

Ground processing begins with raw data archiving followed by decoding and separation of the engineering and science data packets. Data quality and completeness checks flag any missing or corrupt data for retransmission at the next opportunity. Engineering telemetry is channelized and stored in a telemetry database for future query. Science data is decoded, calibrated, time-tagged and archived. Science data post processing is conducted on a monthly cadence to produce calibrated light curves for each target star and quarterly to search those light curves for repeating transit patterns. Data quality assessment, statistical confidence checking and false positive rejection algorithms ensure that the highest quality transit candidates are given top priority when forwarded to the science team for further evaluation.

#### *The Kepler Technique*

In the brief history of planet detection, four techniques have been developed and proven to detect extrasolar planets. 1) The *Radial Velocity* technique measures the spectral Doppler shift that occurs when a star is pulled toward and away from the Earth by its planet(s) [1]. 2) The *Transit* technique measures the minute change in a star's luminosity when a planet orbiting that star moves into the line of sight between



**Figure 2 – An illustration of the transit technique [4]**

the star and Earth. 3) The *Gravitational Lensing* technique looks for slight variation in distortion of a distant object that occur as a planet orbits a lensing star [2]. 4) Finally, the *Direct Imaging* technique, first accomplished in imaging Fomalhaut b [3], images exoplanets directly by attenuating light from the host star to reveal orbiting planets.

The Kepler mission uses the transit technique to detect minute changes in a star's luminosity as an orbiting planet passes between the star and an observer. To reliably detect extra-solar planets using the transit method, a large population of stars must be monitored nearly continuously for several years with a very high precision photometer

capable of detecting less than 40 ppm variation in a star's luminosity.

The large population is required due to fact that the transit technique requires planetary systems whose orbit plane is nearly edge on with respect to Earth. Only planets that actually pass across the disk of their host star, as viewed from Earth, will cause the minute reduction in stellar luminosity that reveals a planet's existence. Assuming random orientations of planetary systems across the galaxy, only 0.5% of planetary systems will be correctly aligned to view transits from Earth.

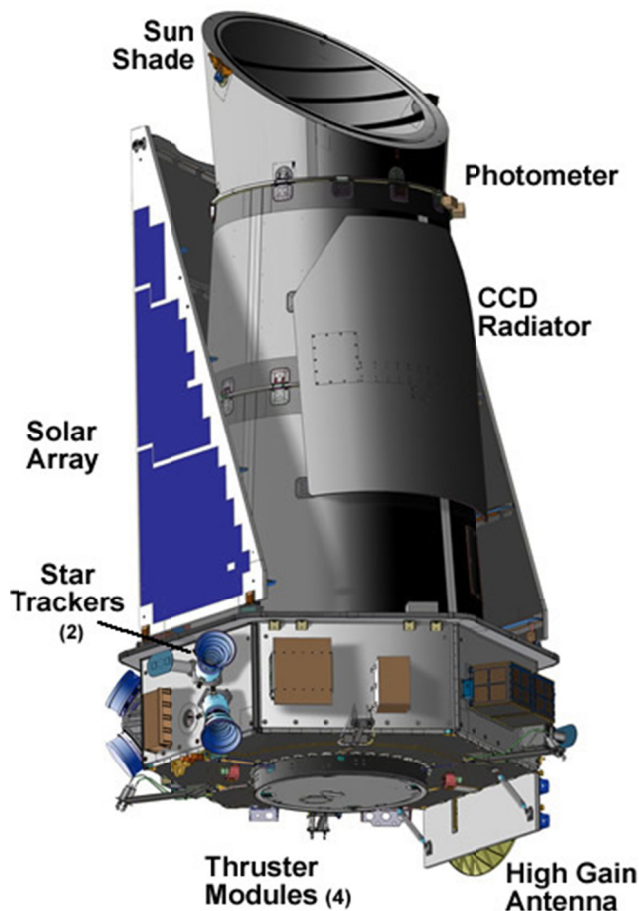
Observations must be performed nearly continuously for extended periods of time because although planetary transits have distinctive signatures (as in Borucki et. al. 2010 [5]), at least two transits are required to allow prediction of a third transit. A fourth transit is often required to eliminate false positives and periodic systemic noise sources. So for an Earth like planet orbiting a sun like star at 1 AU, four transit observations will take four years. Luckily, a great many interesting objects orbit at much shorter periods allowing Kepler to make many discoveries from the 2+ years of data collected to date.

High precision photometry is required due to the large disparity in radius between a star and the terrestrial sized planets Kepler is designed to detect. For example, the transit of the Earth across our Sun would only cause an 84 ppm reduction in the total luminosity of the Sun to a distance observer. Kepler seeks to detect Earth and even Mars sized planets transiting their host star. This requires significant photometric precision and stability.

#### *Flight System Overview*

The Kepler flight system consists of two major components: the photometer and the spacecraft (Figure 3).

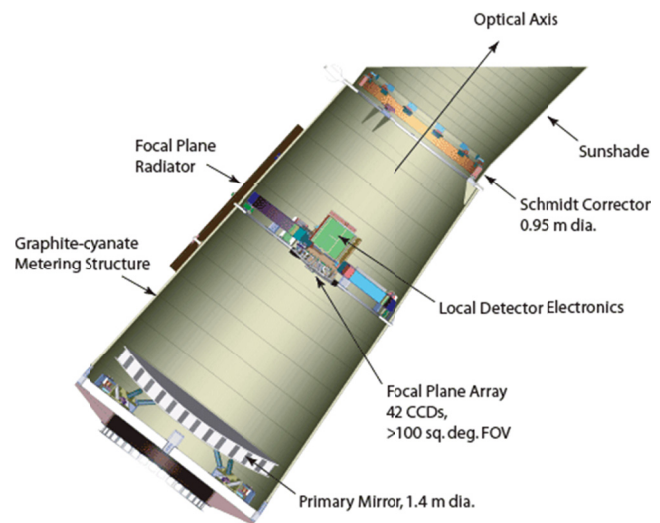
*Kepler Photometer* – The photometer is Kepler's only science instrument and consists of a 9m tall carbon fiber tubular metering structure with a 55° solar shade protecting the 0.95m entrance aperture. Internally, the photometer consists of a 1.4m primary mirror at the base of the metering structure, a 95M pixel CCD array supported at the primary mirror focus by a spider structural assembly and a 7.6 cm thick Schmitt corrector at the entrance pupil with a 0.95m aperture stop. The 95M pixel focal plane consists of 42 individual 1k x 2k CCDs arranged in a 4 fold symmetric pattern. Just behind the CCD array sits the Focal Plane Electronics which provide clocking and high speed readout of the CCDs. The spider assembly, holding the focal plane and focal plane electronics at the primary mirror focus, also supports cabling and a heat pipe assembly. The heat pipe conducts heat away from the focal plane electronics and provides a heat sink for fine thermal control of the focal plane. In addition to the 42 CCD mounted to the focal plane, four fine guidance sensors are mounted at the corners of the



**Figure 3 – Kepler Flight System [4]**

array. These fine guidance sensors are a critical component of the attitude control subsystem that is responsible for keeping pointing jitter below 0.003 arcsec / 15 min (1 sigma per axis) while collecting science data.

*Kepler Spacecraft* – The Kepler spacecraft hosts all of the engineering functions required to support photometer operations. This includes the power, command and data handling subsystems, the attitude control subsystem and the telecommunication subsystem. These engineering subsystems are fully redundant, single fault tolerant and capable of autonomously maneuvering the flight system to a power positive, sun safe attitude in the event of an anomaly. The telecommunications system is comprised of an X band up, X / Ka down transceiver capable of a maximum downlink rate of 4.33 Mbps through a 0.85m high gain antenna. The body fixed solar array provides ample power but requires that the spacecraft be seasonally rotated 90° to maintain sun on the array. Stellar reference is provided by a pair of Ball CT-633 star trackers and inertial reference is provided by redundant Northrop Grumman LN-200S Inertial Reference Units. However, these units are only used for large maneuvers such as safing and monthly data downlinks. The body fixed solar array would provide



**Figure 4 – Kepler Photometer [4]**

significantly more power than required for nominal operations in ideal lighting conditions but is frequently used angled away from the sun by as much as 30°, necessitating the oversized array.

#### *Ground System Overview*

The Kepler ground system consists of the facilities, people, procedures, hardware and software used to operate and maintain the Kepler flight system and to capture, process, archive and distribute data returned from the spacecraft. Kepler ground system facilities include the Deep Space Network (DSN), a Mission Operations Center (MOC) at the Laboratory for Atmospheric and Space Physics (LASP) at the University of Colorado, the Flight Planning Center (FPC) at Ball Aerospace and Technologies Corporation (BATC) in Boulder, the Data Management Center (DMC) at the Space Telescope Science Institute (STScI) in Baltimore, and the Science Operations Center (SOC) at Ames Research Center, Moffet Field, CA. These facilities undertake the accounting, processing and analysis of data; the formulation and transmission of commands to the Kepler spacecraft, and the processing of validated tracking data for orbit determination. Connections between the MOC, FPC, DMC, and SOC use NASA communications network infrastructure.

The DSN provides the spacecraft interface with the ground system. This interface consists of X-band commanding, X-band real time telemetry, Ka-band playback telemetry, and radiometric tracking data. Tracking data is generated for each pass and delivered to the Navigation Team at the Jet Propulsion Laboratory, Pasadena, California which supports maneuver planning and orbit determination. All other transmissions collected from the spacecraft are forwarded to the MOC for further processing, analysis and distribution.

Command, telemetry and science data files are transferred between the DSMS and the MOC using DSN standard protocols and Space Link Extension (SLE). The MOC generates command products, provides DSN contact

planning, processes and monitors engineering telemetry, extracts science data, detects and reports spacecraft anomalies, determines gaps in data from the spacecraft, makes data recall requests, and performs other routine mission operations tasks.

The FPC is the operations support center for the Kepler spacecraft and is responsible for mission operations management, flight segment (FS) maintenance, and spacecraft performance assessment. The FPC receives near real-time engineering telemetry from the MOC in support of critical activities. FPC functions include FS maintenance activities, photometer performance assessment, system test bed operations, command generation & validation for special FS activities, long-term telecom link analyses, FS anomaly resolution, spacecraft flight software maintenance, command & telemetry database updates, and FS status reporting.

The DMC is responsible for processing photometer telemetry and ancillary data received from the MOC into cadence science data sets for further analysis at the SOC. Processing consists of un-compressing science telemetry, converting science telemetry into standard formats, extracting ancillary engineering data, and correcting for velocity aberration & barycentric time. The DMC also calibrates Full Frame Images (FFIs), removing pixel level systematic errors, correcting bias (black-level), smear, and dark-current with collateral data, and applies pixel calibrations from focal plane characterization models. It also archives processed science data from the SOC and manages the science data archive.

The SOC operates the science data pipeline, which is used to analyze stellar photometric data from the Kepler spacecraft for transit candidates. The SOC performs target observation planning, target list maintenance, generation of target and compression tables, detailed photometer performance monitoring and analysis of stellar photometric data. It performs the transiting planet searches and other science data analysis activities that the Kepler science team ultimately uses to identify extra-solar planets. The SOC provides the DMC with data processing pipeline inputs such as target definition, aperture definition, and bad pixel tables. Following receipt of original and calibrated science data set files from the DMC, the SOC returns photometric light curves and target database tables to the DMC for inclusion in the Kepler Data Archive.

## **2. END TO END DATA PATH**

### *Photometer Optical Path*

Light from the science field entering the telescope aperture is reflected off of the primary mirror and focused on the focal plane. Aberrations inherent to the Schmitt telescope configuration are corrected at the entrance aperture by a 7.6cm thick, 0.95m diameter fused silica Schmitt corrector. The focal plane is designed to accommodate the curved nature of the best focus field presented by the telescope.

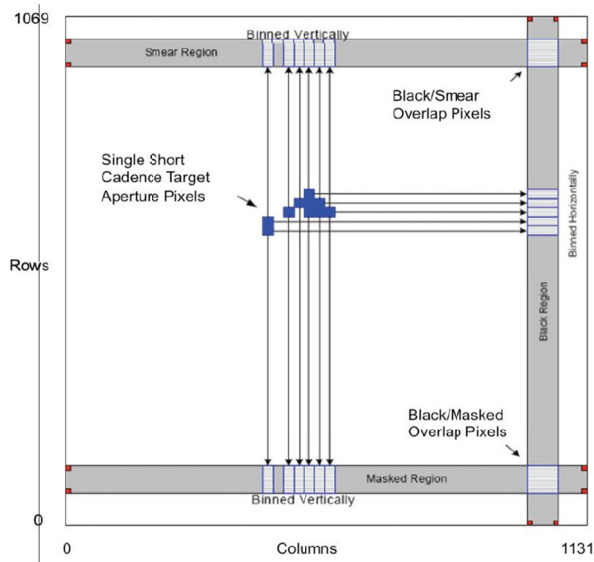
However, since the technology to cost effectively produce curved CCDs does not yet exist, sapphire field flattening lenses placed in front of every CCD perform local correction just before light hits the array.

### *Data collection, processing and compression technique*

Photons striking the focal plane are collected on one of the 42 charge-coupled devices (CCDs) that comprise Kepler's 94.6 megapixel Focal Plane Array (FPA). The CCDs are full frame, 2200x1044 pixel line-transfer devices. A readout amplifier is located on both ends of the serial register and during normal operation one half of each CCD is read out of each output. The 20 rows closest to the serial register are covered with an opaque light shield, so the photosensitive area of the science CCD is only 2200 columns by 1024 rows. Photons are collected in a pixel for an adjustable integration time before being read out in response to a clocking signal from the Local Detector Electronics (LDE). Pixels are clocked one row at a time along the parallel registers into the serial register, then one pixel at a time out of the serial register and through the output amplifiers. Given the large number of pixels and short time available for readout, five high speed output channels are used to process the focal plane. The analog signals from the CCD outputs are digitized, and then co-added for 270 readouts (approximately 30 minutes) by Science Data Accumulators (SDAs). Co-adding serves to amplify the signal and reduce the data stream from the FPA. The data volume is further reduced by only accumulating and saving pixels of interest that are specifically requested by the science team. Target and Aperture tables define the center of a target area, and the number of pixels of interest distributed around each target that should be saved. Reference Pixels (RP) for monitoring photometer performance and Background Targets (BG) for correcting residual instrument artifacts in the image, are also collected. Collateral data are pixels collected from masked smear, virtual smear and trailing black regions at the edges of the CCD (Figure 6). They are used for one-dimensional dynamic correction during science pipeline post processing. Full Frame Images (FFIs) consist of all of the pixels in the FPA; these are collected to verify star location, detector properties, and to characterize image artifacts.

Target apertures are generally a ~13x13 pixel box, with the average target containing 32 pixels of interest – but they do not necessarily need to be rectangular or contiguous. There can be up to 1024 different apertures in total for science and BG targets. There are about 5.44 million pixels of interest (POI) available for science targets, most of which are used for planetary transit search. In parallel to science target data collection, a set of 512 POIs are taken from the image every nine readouts, to provide one-minute short cadence data. This high rate data is used for stellar seismology [8]. Background pixels have their own target and aperture definition tables. Reference Pixels are stored in engineering telemetry and used to monitor photometer health by





**Figure 6 – Kepler CCD Readout Topography [6]**

providing data at more frequent intervals than the downlink of science and background data.

Pixel data collection is managed by the Focal Plane Array (FPA) Data Manager, running on a RAD750 single board computer. The FPA Manager receives interrupt signals to begin collection of science data and the necessary collateral data. Upon receiving an interrupt, the FPA Manager checks the appropriate Target Definition Tables and serially loads the pixels of interest; all other pixels are discarded.

Even with the selective capture of pixels, data volume is still very high and does not meet the constraints imposed by the **230 MB** Solid State Recorder (SSR) and the **4.33 Mb/s** spacecraft downlink bandwidth. In order to keep science data volume within the capacity of the SSR, data is compressed using 3 techniques: *Requantization*, *Baseline Subtraction*, and *Huffman encoding*.

*Requantization* consists of mapping one or more consecutive source quantized values on to the same requantized integer value. In this way, the requantized data range has fewer values than the source quantized data range, and can be expressed by fewer bits. *Baseline subtraction* consists of collecting an uncompressed baseline exposure then subtracting the next 48 exposures from this baseline. Only the difference values (typically values close to zero requiring fewer bits) are stored. *Huffman encoding* is a form of lossless data compression where a symbol string in the source data set is substituted for a code. The most common source symbols are expressed using shorter codes than those used for less common source symbols, resulting in an encoded data set significantly shorter than the source. Requantization, baseline differencing and Huffman encoding of science data reduces the bits-per-pixel from 23 to about 12. This compressed science data is then formed into science source packets and written to the SSR for storage. The spacecraft keeps a correlation table that indexes

the science data against its location on the SSR. Science data is transmitted at Ka-band every 30 days at data rates of up to 4.33 Mb/s. Reference pixels are not compressed and are downlinked every 4 days via X-band at rates up to 16kb/s.

#### *Data transmission operation*

Science source packets from the SSR are wrapped to form Virtual Channel Data Units (VCDUs). Each VCDU is then Reed-Solomon encoded, randomized and prefixed with an Attached Synch Marker (ASM) to form a Channel Access Data Unit (CADU). The CADUs are convolutionally encoded and transmitted. The spacecraft telemetry system downlinks science and engineering data to the Ground Segment via two parallel telecommunication systems. The X-Band system is used to downlink real-time engineering data and reference pixel data at a low data rate via low gain patch antennas every 4 days and the Ka-Band system is used to downlink stored science and stored engineering data at a high data rate via a body pointed high gain antenna every 30 days. The DSN receives CADUs and removes the convolutional, synch word, randomization, and Reed-Solomon encoding. Any CADUs with unrecoverable errors are discarded. The VCDUs containing the telemetry wrapper and CCSDS<sup>1</sup> source packets are extracted and delivered from the DSN to the MOC via Space Link Extension (SLE). Recorded engineering data is downlinked in the same way via different virtual channels (VCs). In case of interrupted data transfer from the DSN to the MOC, the MOC can issue a recall request for data still stored at the DSN. In the event that spacecraft to DSN downlink is interrupted – for example in the case of bad weather – the missing data can be retransmitted from the spacecraft at the next monthly contact. The storage correlation table that maps science data sets stored on the SSR is always downlinked at the beginning of a Ka-band contact. Analysis of the science data and the storage correlation table enables the MOC to identify data gaps and missing packets for possible retransmission. The MOC converts the missing packet information into retransmission commands that specify the SSR locations where playback should begin and the amount of data to retransmit from each location.

#### *Raw Data Receipt and Archiving*

All downlink data successfully received and decoded by the DSN station is passed to the SLE gateway where it is separated into virtual channel streams and stored for retrieval by the MOC. The MOC retrieves the data from the DSN and begins processing the data. The first step in the MOC process is to archive the raw data. Next, the data is reassembled into CCSDS source packets from VCDUs. Stored Engineering packets are then sent to the engineering telemetry processors where they are demodulated, stored in the MOC engineering database and made available to the engineering and operations teams via database query.

<sup>1</sup> Consultative Committee on Space Data Systems (CCSDS)

Science Data packets are reassembled into cadence files. The MOC checks for cadence file completeness by using the SSR Storage Correlation Table to perform an accounting of all science data packets in each cadence file. If there is data missing, retransmit commands are uplinked to the spacecraft at the next monthly contact. Cadence files determined to be complete are packaged and forwarded to the DMC.

#### *Post Processing*

Once received at the DMC, raw cadence files are first archived and then decompressed to form raw pixel files. Raw pixel files are merged and packaged with Ancillary Engineering Data into FITS (**Flexible Image Transport System**) files, archived again then transferred to the SOC. Once the data arrives at the SOC, it is read, validated, and stored locally for further data processing. There are two main pipelines that are run in the SOC. The first, which is run following the monthly Science Data download, produces calibrated light curves. The other, which is run following the completion of an observing quarter, performs the transit search and data validation.

The Monthly SOC pipeline begins with pixel-level calibration (CAL) to produce calibrated target and background pixels prior to further photometric analysis. CAL performs basic corrections to remove instrument signatures. Many of the corrections are typical of photometry (black correction, gain and non-linearity correction, flat field correction), but some are unique to Kepler, including smear and dark current correction due to the lack of a shutter on the Kepler photometer [7].

Following CAL, the calibrated pixels are processed by the Photometric Analysis (PA) module. PA produces raw light curves by removing sky background and performing simple aperture photometry by summing the calibrated corrected signal from each pixel of interest in an individual target aperture. PA also measures the centroid of each target in each cadence measurement.

The final step in the monthly production of light curves is the Pre-Search Data Conditioning (PDC). PDC removes systematic errors from sources as pointing drift, focus changes, and thermal transients. The outputs of the monthly processing pipeline include raw and calibrated pixels, raw and corrected flux time series, centroids, and other statistical parameters for each target star. Following the completion of an observing quarter, these products are returned to the DMC for archive and eventual public distribution.

#### *Candidate identification and selection*

Planetary candidates are produced by the Quarterly processing pipeline, which consists of two modules – the Transiting Planet Search (TPS) and Data Validation (DV).

As the name implies, the quarterly pipeline is run once an observing quarter is complete and all of the data from the

quarter have been returned to the ground. The first step in the quarterly pipeline is Transiting Planet Search (TPS). TPS utilizes a wavelet based adaptive matched filter to look for transit like features in the light curves. Transit like features whose detection statistic exceeds 7.1-sigma are designated as Threshold Crossing Events (TCE), and are passed on to Data Validation for further analysis [7].

DV performs a series of statistical tests to evaluate the confidence in the detection, to reject false positives due to background eclipsing binaries, and to extract physical parameters of each system. Once a TCE has been processed by DV, its signature is removed from the light curve, and the light curve is searched for additional transits. This process (TPS-DV) continues until no further TCEs are identified.

Following the completion of the quarterly pipeline, TCEs are forwarded to the TCE Review Team (TCERT) for further analysis. The TCERT evaluates each TCE utilizing a number of tools and the most promising candidates are forwarded to the Follow-up Observing Program (FOP) for confirmation.

#### *Data Quality Assessment*

There are numerous Data Quality checks built into the SOC processing, both automatic and manual. Automatic data quality checks are focused on measuring the performance of the photometer, and therefore the quality of the science data. The Photometer Data Quality (PDQ) and Photometer Performance Assessment (PPA) pipeline modules perform these automatic data quality checks. The PDQ module operates on the reference pixels that are downlinked from the Flight Segment during the semi-weekly X-band contacts. The PPA module performs more detailed performance and quality checks on the light curves and pixel data associated with all targets.

There are also numerous manual quality checks performed during the SOC pipeline processing. These range from an analysis of operational events that occurred on the spacecraft and flagging the occurrence of these events in the science data, to detailed examination of the data products produced by the SOC pipeline by members of the SOC and Science Office teams.

### **3. KEPLER OPERATIONS**

#### *Nominal Mission Operations*

Kepler mission operations activities execute on three basic cadences: weekly, monthly, and quarterly. Weekly operations activities include mission planning, routine housekeeping, command upload and health & safety monitoring. Monthly operations activities include stored engineering and science data download, flight segment maintenance and monthly data processing. Quarterly operations activities include target management, flight segment re-orientation and quarterly data processing.

### *Weekly Operations Cadence*

A typical operations week for the Kepler team is focused on the twice-weekly DSN contacts with the spacecraft. During these 8-hour X-band contacts, the Operations Team at the MOC performs routine housekeeping activities, Flight Segment health and safety assessments, and command uploads. Telemetry downloads include various flight software log files, low-rate stored engineering data, the SSR Storage Correlation Table, and Reference Pixel files. Additionally, the real-time telemetry stream is checked for anomalous readings or flight rule violations. Command uploads typically include a small number of maintenance commands, clean-up commands, and command sequences required to prepare for an upcoming monthly or quarterly contact.

Following the completion of a semi-weekly contact, downlinked engineering telemetry is processed and archived for use by the engineering team at the MOC and FPC, as well as the larger Mission Operations Team. Reference pixel files are packaged by the MOC and sent to the SOC for processing and analysis. The SOC and Science Office analyze the reference pixel files to perform periodic photometer health assessments and trending as well as quick look Flight Segment pointing performance assessment.

Once a week, the Mission Operations Team gathers for an Operations Status and Planning Meeting. At this meeting, the health and status of the entire Kepler Mission System (Flight Segment, Ground Segment, and Photometer) are reviewed. The operations team also reviews upcoming activities and changes prior to approval for command product generation, validation, and upload.

### *Monthly Operations Cadence*

Once per month, Kepler suspends science data collection, turns its high gain antenna towards earth and downlinks stored engineering and Science Data via the high rate Ka-band system. These activities, occur during a 24-hour DSN contact, and typically result in a 15-16 hour break in science data collection.

Monthly contact activities actually begin prior to the initiation of the Ka-band DSN contact. Initial monthly activities include the stopping science data collection, collection of monthly calibration data (reverse clock data and a full field image), turn-on of the Ka-band telemetry stream and maneuvering the high gain antenna to Earth point. Once the Ka-band telemetry link is established, on-board sequences command the downlink of high-rate stored engineering data, the current month's new science data, and any data from the previous month that required retransmission. Downlink of a full month of science data at Ka-band downlink rates usually takes between 4 and 6 hours, depending on the volume of retransmission data.

Once the Ka-band downlink is complete and the Ka-band link is turned off, the flight segment is maneuvered back to

the science attitude. This maneuver is followed by a period of thermal settling to allow the star trackers and fine guidance sensors to return to alignment. Once the thermal alignment period is complete, photometer fine pointing is reestablished and science data collection is resumed.

In addition to the science data downlink, normal semi-weekly housekeeping activities are conducted as well as uplink of new short cadence target tables and a new background command sequence.

### *Quarterly Operations Cadence*

Every three months (once per quarter), the Flight Segment must be re-oriented so that the solar arrays point back at the sun and the CCD radiator is repointed at deep space. This re-orientation requires that the spacecraft be rolled 90 degrees about the photometer bore sight. In order to minimize the impact to data collection, these re-orientations, or "Quarterly Rolls" are performed in conjunction with a monthly science data download.

Quarterly roll contacts begin in the same manner as the monthly contacts, with the stopping of science data collection, calibration data acquisition, maneuver HGA to Earth Point and Ka-band science data downlink. The differences come following the completion of Ka-band downlink. At this point, instead of returning to the previous science attitude, the Flight Segment is rolled 90 degrees to the next season's science attitude. Following the roll, the Flight Segment enters a period of thermal stabilization and alignment. During the thermal stabilization period, the operations team ensures that the target tables for the upcoming quarter of science data collection are loaded and ready for use. The Operations Team also ensures that the new quarter's fine guidance sensor pixel of interest table is readied. Following completion of the thermal stabilization period, the Flight Segment is returned to fine point mode in the new orientation and initial calibration photometer data is collected. A typical quarterly roll contact is scheduled for 36-48 hours of DSN coverage and results in 22-24 hour science break.

Planning for the quarterly roll actually begins almost 90 days prior to the execution of the roll with the initiation of the Target Management Process. Since the focal plane of the photometer is re-oriented 90 degrees, the target stars will fall on new locations on the CCDs. As a result, new target tables must be generated and uplinked to the Flight Segment. This processes, executed by the Science Office and the SOC, takes target requests from the Science Team and the Guest Observer Program and generates the list of targets that will be observed in the upcoming quarter. The SOC then prepares the target tables needed to collect the required Science Data.

### *Anomaly Resolution and Recovery*

Kepler is a data-driven mission, and maximizing data collection time is a key operating parameter. As a result, the

primary focus of the Kepler Anomaly Process is the safe return to science data collection following an on-board anomaly.

The Kepler spacecraft provides onboard fault detection and response functions geared mainly toward guaranteeing the safety of the spacecraft. Most onboard fault conditions result in a system safing event where the spacecraft turns off the photometer, slews to point a vector normal to the body fixed solar array directly at the sun and begins a slow 0.1 RPM roll about the solar array normal vector. This places the spacecraft in a stable, thermally safe orientation that guarantees sunlight will not enter the photometer boresight. Telecommunication drops to low rate through the -z low gain antenna. The spacecraft will wait in this configuration for ground intervention.

The Kepler Anomaly Response Process is executed by the Project Systems Engineer (PSE) on behalf of the Project Manager (PM). The process begins when the Mission Operations Team at the MOC recognizing that an anomaly has occurred. Appropriate members of the Flight Segment Engineering and Operations Management Teams are notified and the Anomaly Response Process is initiated. The Engineering Team's prime focus during the early stages of an anomaly response is to assess the health and status of the Flight Segment and to begin to formulate an understanding of what anomaly has occurred.

Shortly after the initiation of the anomaly response, the PSE convenes an Anomaly Review Board (ARB) consisting of Engineering, Science, Operations, and Management personnel. The ARB's primary role is to evaluate information pertaining to the anomaly and Flight Segment health and authorize the execution of diagnostic and recovery actions. The ARB may choose to form an Anomaly Resolution Team (ART) consisting of subject matter experts to further diagnose the problem and to develop recovery and/or mitigation actions. Once the ARB is satisfied that appropriate recovery actions are available to allow a safe return to Science Operations, the ARB returns control of recovery and operations activities to the Kepler Mission Director (KMD).

ART activities may continue following the return to science operations. These activities may include root cause assessment, operational risk assessment, and development of Flight Segment and operational changes to mitigate issues presented by the anomaly. Once the ART activities are complete, their results are presented to the ARB for implementation approval. The Operations Team is then responsible for implementation of these changes under the direction of the KMD. Finally, the ARB reviews the results of implementation, and if successful, approves closure of the anomaly report.

#### 4. SUMMARY

Teasing out the telltale signature of a planetary transit from raw, noising photometric data requires a careful and well-

coordinated data management pipeline. The relatively smooth end to end operation of the Kepler data pipeline has produced 25 confirmed planet discoveries with over 1200 candidates currently identified and awaiting independent confirmation. Identifying planetary candidates through the transit method, once dismissed as impossibly difficult, has proven robust. The discoveries that Kepler is currently making are sure to have resounding effects on planetary science, astrobiology and cosmology.

#### REFERENCES

- [1] Campbell, B. ; Walker, G. A. H. ; Yang, S. (1988). "A search for substellar companions to solar-type stars". *Astrophysical Journal* **331**: 902–921.
- [2] Mao, S. & Paczynski, B. (1991) "Gravitational microlensing by double stars and planetary systems". *Astrophysical Journal* **347**: L37
- [3] Kalas, P. et. al. (2008). "Optical Images of an Exosolar Planet 25 Light-Years from Earth". *Science* Vol. 322 no. 5906 pp. 1345-1348
- [4] Kepler Public Website: [kepler.nasa.gov/](http://kepler.nasa.gov/)
- [5] Borucki, et. al. (2010). "Kepler-4b: A Hot Neptune-Like Planet Of A G0 Star Near Main-Sequence Turnoff". *Astrophysical Journal Letters* **713**: L126
- [6] Kepler Instrument Handbook, Astrophysics with Kepler website: <http://keplergo.arc.nasa.gov/calibration/KSCI-19033-001.pdf>
- [7] Jenkins et. al. (2010). "Overview of the Kepler Science Processing Pipeline". *Astrophysical Journal Letters* **713**: L87
- [8] Christensen-Dalsgaard, et. al. (2006). "Asteroseismology with the Kepler Mission". *Comm. In Asteroseismology* Vol 150



## Acknowledgements

The work presented was carried out by the Kepler Project as part of the NASA Discovery Program. The Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, under a contract with the National Aeronautics and Space Administration, was responsible for management of Kepler's development and launch. NASA Ames Research Center, Moffett Field, CA manages Kepler mission operations and ground data system. Ball Aerospace & Technologies Corp. of Boulder, Colo., was responsible for developing the Kepler flight system and supporting mission operations.

Approved for unlimited release: JPL CL#11-4970

## Biographies



**Brian Cooke** received a B.S. in Engineering Science and Mechanics from Virginia Tech in 1995. He has been with JPL for more than 15 years and is the recipient of three NASA Exceptional Achievement Medals. He is currently the Project System Engineer for the Europa

Habitability Mission helping to develop and plan NASA's exploration of this intriguing Jovian moon. He has previously served as the Kepler Project System Engineer, Dawn Project V&V Engineer and the GALEX Instrument I&T Manager.



**Richard Thompson** received a B.S. in Engineering in 1986 and an M.S. in Operations Research in 1987 both from the University of California, Berkeley. He has been with Orbital Sciences Corp. for 10 years. Most recently, he served as the Project Systems Engineer for the Kepler Mission. Prior to that, he was the Ground Segment Manager and

Ground Segment Systems Engineer for Kepler.



**Shaun Standley** is the Group Supervisor for the System Verification and Validation Group at JPL. Before this he held a number of positions in V&V Engineering and spacecraft operations, including Project V&V Engineer for the Soil Moisture Active and Passive (SMAP) mission, End to End Information System engineer for the Kepler mission, Spacecraft Engineer for the Cassini orbiter, Systems Engineer for the Huygens Titan probe, and Systems Engineer for the Ulysses Solar orbiter. He has a B.Sc. (Hons) from Southampton University in the UK where he also did postgraduate work, and an MBA from Pepperdine University